

Making the Invisible Visible in Writing Classrooms: An Approach to Increasing Textual Awareness using Computer-Aided Rhetorical Analysis

Necia Werner, Carnegie Mellon University

Dr. Necia Werner is an Assistant Teaching Professor of English and Director of the professional and technical writing programs at Carnegie Mellon University. Werner serves on the advisory committee (AdCom) of the IEEE Professional Communication Society, and as an Associate Editor for the teaching case section of the IEEE Transactions on Professional Communication.

Suguru Ishizaki, Carnegie Mellon University

Suguru Ishizaki is an Professor of English in the Department of English at Carnegie Mellon. His current research interests include pedagogy of communication and design for students and professionals in the technology/engineering disciplines, and computer-aided rhetorical analysis.

Making the Invisible Visible in Writing Classrooms: An Approach to Increasing Textual Awareness using Computer-Aided Rhetorical Analysis

Introduction

Writing requires countless composing decisions that are typically beyond the writer's conscious grasp. For students, writing can feel like a process that they have little control over, and a skill that only a certain few possess. Much of the skill in being "text-aware" involves understanding that texts produced from classroom assignments are not just composed of words and sentences, but of highly structured and often highly predictive composing decisions. A fundamental goal of core writing courses in many first-year writing and upper-level technical writing services courses for STEM majors is to impart this textual awareness to students, helping them understand that different decision-making at the compositional level leads to different text types appropriate for specific purposes and audiences. However, "visualizing" the decision-making processes of underlying writing, and "seeing" textual patterns within genres, is an extremely abstract idea for students, and even harder to teach within the constraints of a single course. It's a bit like investigating dark matter; one can hypothesize without seeing it that it exists because of radiation signals observed when particles collide, but not really understand what it's composed of or how it holds galaxies in the universe together. Our goal in this project is to provide student writers with a means to jumpstart their understanding of writing as compositional decision-making by equipping them with the means to quickly, and literally, "see" their composing decisions.

To help students notice and reflect on composition decisions in their writing, our team of writing researchers, educators, and statisticians is piloting use of a suite of computer-aided learning tools for corpus-based text analysis in core writing courses. Developed and refined over the last 20 years, the DocuScope tool has heretofore successfully demonstrated its strength as a research tool to sort corpora into identifiable genres, for example, identifying the statistically significant patterns and moves that differentiate histories, comedies, and tragedies in Shakespeare's plays [1, 2], as well as its potential as an educational tool in writing courses. At Carnegie Mellon University, the tool has been used for these purposes in a graduate-level writing course for design students [3, 4], which created a writing classroom environment that functioned like a critique-based design studio; a corpus of student texts from the class could be analyzed in aggregate to visualize the rhetorical patterns common to a genre [5], and students could see where particular elements in their own draft differed most noticeably from those of classmates. From the instructor perspective, "seeing" was learning (see Figures 1 and 2), in the sense that students could see that genres are composed of predictive patterns, but that individual writers have choice as well. The result is a classroom in which writing can be discussed as an act of representational composition, or textual brushstroke choices made consciously in the service of creating a particular reading experience for readers [6].

The goal of the current project is to scale up this approach to support core, undergraduate writing courses with multiple sections and instructors, starting with courses in our first-year writing program and technical communication program for STEM majors. These courses reach the most undergraduates on our campus and, ideally, learning outcomes are scaffolded meaningfully from one to the next in ways that are pedagogically sound and can be communicated to students.

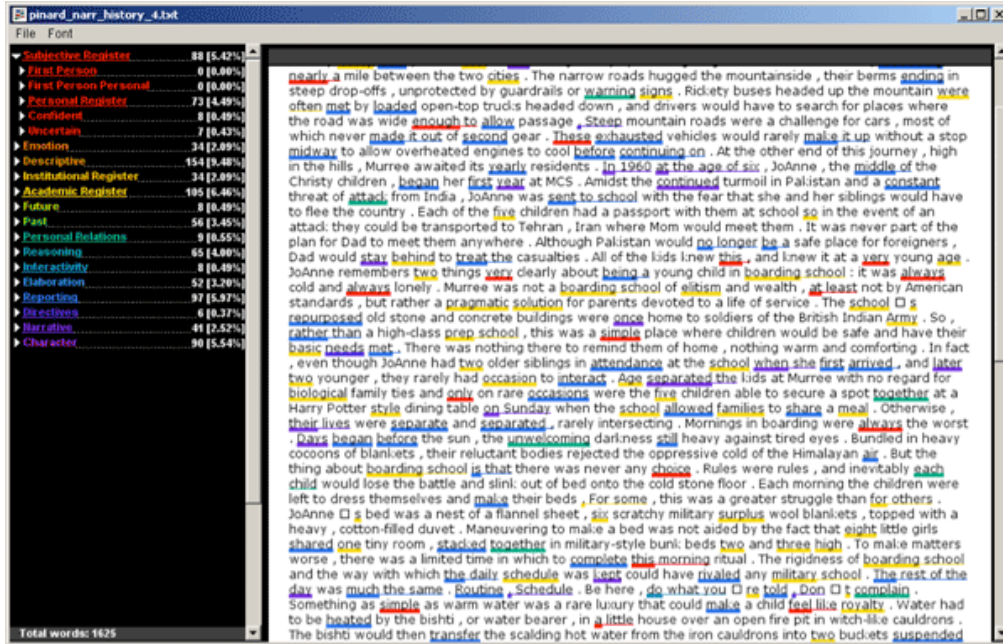


Figure 1. Screenshot of individual text analyzed and marked for textual patterns.



Figure 2. Screenshot of analysis of a corpus of multiple genres to see patterns of similarity and different across genres, as well as individual outliers within the corpus.

A Tool for Helping Instructors Analyze Assignment Genres

At many institutions, including our own, the first-year writing (FYW) program and subsequent technical communications (TC) service course reach the largest number of students with writing instruction. For many undergraduates, these courses comprise the entirety of their writing-focused coursework. Ideally, then, these foundational courses are meaningfully linked for students, and lay the groundwork for them to notice how their learning can carry forward into other classes and genres. At our own institution, the need to articulate learning outcomes across these core courses recently came to the fore; in response to updates to the general education and core requirements of our STEM units, particularly those of computer science and our largest campus unit, engineering, we tripled the annual number of technical writing service course sections to accommodate increased demand.

Although we felt confident that the assignments and learning outcomes for our FYW and TC courses were well aligned, articulating that connection to students and other campus stakeholders is a perennial challenge. How do you *know* that your courses are sequenced effectively? Additional challenges to building what we've dubbed a "transparent corridor" of learning outcomes between classes are many and well documented, including lack of student motivation to learn writing skills, a miscalibrated self-assessment of abilities, and beliefs that writing skills are a "gift" as opposed to something that can be externalized and taught.

In spring 2016, we received funding and IRB approval for a project to investigate the potential for technology-enhanced learning tools to accelerate student learning and build a strong pedagogical bridge between our FYW and TC courses. Our goal for the overall project is to develop a suite of computer-aided learning tools that will help students in all disciplines *notice hidden structures and composing decisions in writing, and become more self-aware and reflective writers*. As a first step in Phase 1 of the project, we investigated whether our writing assignments are, in fact, scaffolded appropriately for students, and whether they build meaningfully upon one another, not just within one course, but across courses. What would a corpus analysis of our students' texts reveal about the most prevalent rhetorical moves students evince in a single assignment? Are some of those same strategies prevalent in subsequent assignments, and if so, which ones? What *new* rhetorical moves are evinced?

In spring 2016, a corpus of student project texts was collected and anonymized from five sections of FYW and eight sections of TC. The FYW courses selected for this analysis included a sequence of three student paper projects: a comparative genre analysis (CGA), a proposal for an original research contribution, and a final contribution paper. For the TC course, the first two assignment genres were collected: a cover letter (part of a résumé and cover letter package for a job application, a common assignment in technical communication service courses) and a proposal directed to a real decision maker (either a proposal to receive funding for a research project or a proposal to a campus decision maker to support an idea for policy change). In FYW, the emphasis is on introducing students to the kinds of writing they can expect in an academic context, with emphasis on information literacy and using sources strategically. In TC, the emphasis shifts from writing in an academic context to writing in professional contexts for largely non-expert audiences.

Given a corpus of texts, the DocuScope tool for corpus rhetorical analysis draws on a roughly 50+ million patterns of English across over 7,000 hierarchically organized categories with zero overlaps [7]. Texts can be analyzed and explored individually, as in Figure 1, where individual categories can be selected or deselected in the text to highlight how particular language choices map onto particular categories (e.g., uses of First Person discourse). Alternatively, texts can be analyzed in aggregate to surface the rhetorical categories that distinguish one genre from another (e.g., Shakespeare’s plays [1, 2]).

Analysis of our initial corpus of texts indicated that students were, in fact, using different rhetorical moves as they moved through the assignments (see Table 1), and that these skills built upon one another in ways that both mapped onto and sharpened our understanding of the kinds of moves we want our students to see in different genres. For example, in FYW, students were learning and practicing moves for citing sources (“Citation”) across all genres, which was picked up again when they were writing decision-maker oriented proposals in TC. It also made sense that the language of comparison (“Compare”: e.g., words and phrases like “same,” “different,” “fewer than,” “less than”) would be highly present in the CGA assignment, as students learn to see and describe genre differences. Interestingly, the granularity of the dictionary categories allowed us to drill down and see when and where students were using *different* kinds citation moves; for example, *authority citations* where the writer signals the validity of a cited source (e.g., “X has shown/demonstrated”), *controversy citations* where the writer signals that a source has not been completely vetted (e.g., “X purports”), and *neutral citations* that do not signal the truth value of the citation either way (e.g., “X states”).

As Table 1 illustrates, a new set of rhetorical categories and genres conventions were lighting up for students as they moved into TC; for example, in the cover letter assignment, students combine strategies for narrating their personal strengths and skills (“Narrate”), and describing positive actions, beliefs, and outcomes (“Positive Values”).

Table 1 Most Frequent Categories of Rhetorical Moves Across Course Assignments

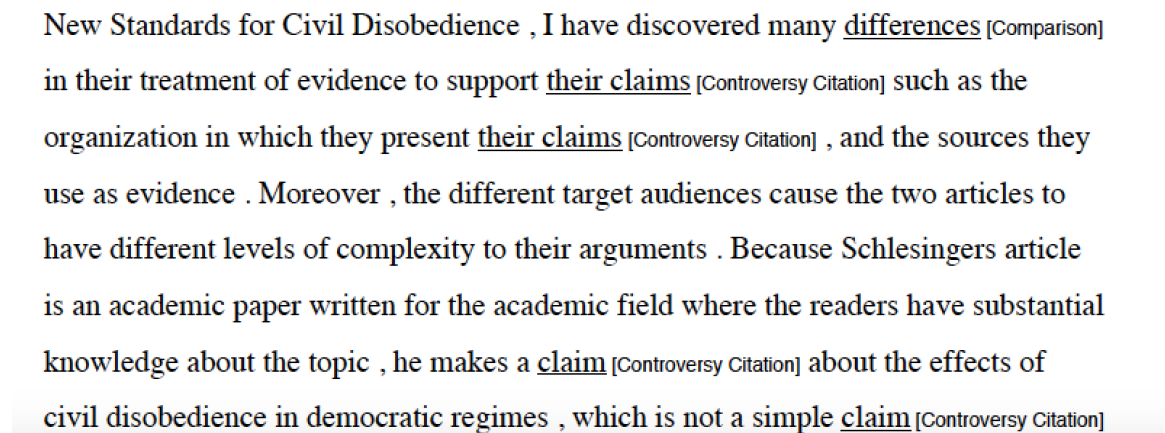
Categories of Rhetorical Moves	First-Year Writing (FYW) Course Assignments			Technical Communication (TC) Course Assignments	
	FYW-CGA	FYW-Contribution	FYW-Proposal	TC-Cover Letter	TC-Proposal
Compare	x				
Citation	x	x	x		x
Proper Names	x				
Personal Pronouns	x				
Public		x	x		x
Inquiry		x	x		x
Swales Introductions		x	x		x
Swales Gaps			x		x
Reasoning		x	x		x
Interactive				x	
Letter Genre				x	
Positive Values				x	
Facilitate				x	
Autobiography				x	
Narrative				x	x
Future				x	x
Strategy				x	x

A Tool for Helping Students Visualize Writing Patterns

In previous courses that used DocuScope to analyze and compare student texts [3, 4], the process for generating data for students (see Figures 1 and 2) were not fully automated. Running the tool therefore required a manual process that had to be handled by the instructor—the creator of the tool, who was intimately familiar with how to use it and how to interpret the output, and versions of the course were co-taught with a statistician. This original process therefore worked as well as it did because (1) the instructors were experts who were extremely familiar with the tool, and (2) were able to assist students in interpreting the analysis. In order to scale the use of this tool for core writing courses, our team prioritized creating new user-friendly tools, capable of presenting the DocuScope results clearly to non-writing experts (i.e., students), as well as to non-experts in corpus rhetorical analysis (i.e., many instructors).

As a critical part of our pilot project, we therefore asked the following questions: What are optimal ways to integrate automated reporting into undergraduate writing instructions? How can these reports be integrated meaningfully for students in FYW and TC? To what extent can these reports positively impact student understanding of structures and composition decisions in their own writing? This phase of the project can accurately be described as: “make the results of the tool usable and accessible for non-expert students and instructors.”

In summer 2016 and fall 2017, as part of our efforts in Phase 2 of the project, we developed a second TEL tool to make headway on this goal. First, we developed an automated process for instructors to upload student texts into DocuScope, analyze texts for the expected rhetorical and linguistic patterns for the genre (as derived through analysis from Phase 1), and then produce an individualized report for students that helped them understand what was “going on” in their text. In one portion of the report, students receive a fully marked up copy of their paper, showing where in the text the tool was marking particular categories of information. Anecdotally, students in the FYW pilot shared surprise that different kinds of citation strategies existed, and that in some cases, they were favoring one type (see “claims” in Figure 3) reflexively and without consideration for rhetorical effect.

The image shows a screenshot of a student report. The text is highlighted in a light gray box. The text reads: "New Standards for Civil Disobedience , I have discovered many differences [Comparison] in their treatment of evidence to support their claims [Controversy Citation] such as the organization in which they present their claims [Controversy Citation] , and the sources they use as evidence . Moreover , the different target audiences cause the two articles to have different levels of complexity to their arguments . Because Schlesingers article is an academic paper written for the academic field where the readers have substantial knowledge about the topic , he makes a claim [Controversy Citation] about the effects of civil disobedience in democratic regimes , which is not a simple claim [Controversy Citation]". The words "differences", "their claims", "claim", and "claim" are underlined, and each has a small bracketed label next to it: "[Comparison]", "[Controversy Citation]", "[Controversy Citation]", and "[Controversy Citation]" respectively. The word "claim" appears twice in the text.

New Standards for Civil Disobedience , I have discovered many differences [Comparison] in their treatment of evidence to support their claims [Controversy Citation] such as the organization in which they present their claims [Controversy Citation] , and the sources they use as evidence . Moreover , the different target audiences cause the two articles to have different levels of complexity to their arguments . Because Schlesingers article is an academic paper written for the academic field where the readers have substantial knowledge about the topic , he makes a claim [Controversy Citation] about the effects of civil disobedience in democratic regimes , which is not a simple claim [Controversy Citation]

Figure 3. Screenshot of where citation types appear in an individualized student report.

In a second portion of the individualized student reports, students were given an abbreviated DocuScope analysis that showed students how *they compared to other students in the class* in their use of a particular categorical dimension in their paper. Figure 4 illustrates a subsection of one student's individualized report on a comparative genre analysis (CGA) paper for the FYW course; the black dot represents the individual student within the box plot, and in this case, indicates that the student's paper contained more comparison language than the class as a whole. It is important to note that this output is not an assessment of the students' papers: the goal is not to have students identically reach a median. The goal is instead to *promote reflection*: as in, "it looks like your paper showed more comparison language than the median for your peers' papers...can you talk through your strategies and explain why?"

Comparison


References comparison words like "compare", "comparison", "similar", "resemblance," "same", "different"; comparative terms with "than" like "less than" "fewer than" and "more than".



●=You. On average, 13.60 patterns were used per 1,000 words

Authority Citation

References citation when the writer signals the truth and validity of the cited source.



●=You. On average, 0.80 patterns were used per 1,000 words

Controversy Citation

References citation that signals controversy and that the cited source has not been completely vetted.

Figure 4. Screenshot from an individualized report, showing where a student fell alongside classmates in dimensions related to citation usage.

A Tool for Promoting Student Reflection

To further our goal of using the individualized reports as a springboard for student reflection, statistician colleagues created an additional technology-enhanced learning tool designed specifically to invite student discussion and comparison. Specifically, we wanted a way to identify, given a corpus of texts in a single class, which of the students in the class were *most different* from one another along a particular category. For example, if student X was well below the median on controversial citations, while student Y was well above the median, how would they account for their differences to one another if they were paired, or "matched" during a peer reflection activity?

Figure 5 is a screenshot from a "matchmaking" tool designed to compare students' texts within the class, and suggest best "matches" for conversation. The instructor can use the tool to toggle between peer reflection group size (2, 3, or 4 students), with suggestions for student peer groups visualized in the form of groups of faces. Instructors can additionally toggle between categories to make informed decisions about the kinds of conversations that would be most productive or interesting. Instructors can project this visualization tool to the class if they so choose, helping students "see" how different students are making different choices in their papers, and priming

them for their subsequent reflection conversation. In this way, the reports serve as the basis for reflection and analysis, as well as in-class peer review activities in which students can debate, discuss, and justify their compositional decisions.



Figure 5. Screenshot of a matchmaking tool to pair students in the class according to the dimensions where they differ most.

Conclusion and Future Directions

Fostering students' ability to notice genres as recurring rhetorical and linguistic patterns rather than classroom-specific isolates, and to see their own composing decisions in these genres as decisions that can be articulated, is ironically both a central tenet of many writing classrooms and, arguably, the most challenging to teach. Our project offers a means for giving students new ways to notice text surface alongside genre, and to combine situation with language in their compositional decision-making. We believe that this awareness of text surface will travel with students throughout their university experience and beyond. Additionally, the project also contributes toward a research-based curriculum articulation and alignment for foundational writing instruction at the university level. At our university, the students most affected by this articulation are our STEM majors, for whom the technical communication course fulfills an advanced writing course. Our aim in this project, to be continued and refined further, is to help student writers see that their composing moves are, in fact, decisions that can be adjusted and refined to meet the rhetorical demands of different genres. This pilot is ongoing; at present, we are engaged in Phase 3 of our project, integrating the DocuScope tool and newly created toolkit materials for teachers and students into a subset of both the FYW and TC writing classrooms.

Acknowledgements

The authors would like to thank the anonymous reviewers of this paper for their feedback. This research is supported by a ProSEED/Crosswalk Seed Grant and has received IRB approval from Carnegie Mellon University.

References

- [1] J. Hope and M. Witmore. "The hundredth Psalm to the tune of 'Green Sleeves': Digital approaches to Shakespeare's language of genre." *Shakespeare Quarterly*, vol. 61, no. 3, pp. 357-390, 2010.
- [2] J. Hope and M. Witmore. "The very large textual object: A prosthetic reading of Shakespeare." *Early Modern Literary Studies*, vol. 9, no. 3, pp. 1-36, 2004.
- [3] D. Kaufer, C. Geisler, P. Vlachos and S. Ishizaki, S., "Mining textual knowledge for writing education and research," in *Writing and Digital Media*, L. v. Waes, M. Leijten, and C. Neuwirth, Eds. Oxford, UK: Elsevier Science, 2006, pp. 115-130.
- [4] D. Kaufer, S. Ishizaki, J. Collins and P. Vlachos, "Teaching language awareness in rhetorical choice using iText and visualization in classroom genre assignments." *J. Bus. Tech. Commun.*, vol. 18, no. 3, pp. 361-40, 2004.
- [5] D. Kaufer, S. Ishizaki, B. S. Butler and J. Collins, *The Power of Words: Unveiling the Speaker and Writer's Hidden Craft*. Mahwah, NJ: Lawrence Erlbaum Associates, 2004.
- [6] D. S. Kaufer and B. S. Buter, *Designing Interactive Worlds with Words: Principles of Writing as Representational Composition*. Mahwah, NJ: Lawrence Erlbaum Associates, 2000.
- [7] S. Ishizaki and D. Kaufer, "Computer-aided rhetorical analysis," in *Applied Natural Language Processing: Identification, Investigation and Resolution*, P. McCarthy and C. Boonthum-Denecke, Eds. Hershey, PA: IGI Global, 2011, pp. 276-296.